

Drift-Free Humanoid State Estimation fusing Kinematic, Inertial and LIDAR sensing

Maurice F. Fallon¹, Matthew Antone, Nicholas Roy and Seth Teller²

Abstract—This paper describes an algorithm for the probabilistic fusion of sensor data from a variety of modalities (inertial, kinematic and LIDAR) to produce a single consistent position estimate for a walking humanoid. Of specific interest is our approach for continuous LIDAR-based localization which maintains reliable drift-free alignment to a prior map using a Gaussian Particle Filter. This module can be bootstrapped by constructing the map on-the-fly and performs robustly in a variety of challenging field situations. We also discuss a two-tier estimation hierarchy which preserves registration to this map and other objects in the robot’s vicinity while also contributing to direct low-level control of a Boston Dynamics Atlas robot. Extensive experimental demonstrations illustrate how the approach can enable the humanoid to walk over uneven terrain without stopping (for tens of minutes), which would otherwise not be possible. We characterize the performance of the estimator for each sensor modality and discuss the computational requirements.

I. INTRODUCTION

Dynamic locomotion of legged robotic systems remains an open and challenging research problem whose solution will enable humanoids to perform tasks and reach places inaccessible to wheeled or tracked robots. Several research institutions are developing walking and running robots with a range of form factors, from all-terrain quadrupeds operating outdoors to experimental bipedal runners which have yet to leave the laboratory.

Locomotion is a typical closed-loop control problem whose primary input consists of the state of the robot — namely, the 6-DOF pose and velocity of its pelvis, as well as the configuration of its joints. Accurate, timely estimates of the robot state not only facilitate effective control for dynamic whole-body motions such as walking, but also enable a greater degree of task autonomy via consistent knowledge of the surrounding environment and the locations of objects within it.

A. Related Work

One common class of estimation method is based on dynamics (e.g., [1]), and relies on knowledge of the controller outputs and a motion model of the humanoid to infer the state of the robot’s centers of mass and pressure. Errors in link center-of-mass modeling and the presence of unpredictable




Fig. 1. The Atlas robot contains 28 hydraulically actuated joints, and its primary sensing is provided by the Carnegie Robotics Multisense SL sensor head which is equipped with a rotating LIDAR scanner and a stereo camera. (photo credits: Boston Dynamics and CRL)

forces (e.g., the from the robot’s power/pressure hoses or external pushes) are accounted for by appending an additional process model for each class of disturbance.

In [2], the authors extend this approach and apply it to the Atlas robot (which we are also using in our work). They discuss the computational challenges of formulating a single extended Kalman filter (EKF) for a humanoid with many degrees of freedom, and propose instead to estimate the pelvis position and joint dynamics in separate filters.

An EKF-based estimator is presented in [3] for a quadruped that uses a sensor-based prediction model and creates filter corrections using foothold measurements. This approach incorporates the positions of footholds into the state vector (using a point model for each foot) and gives particular consideration to consistency and observability analysis. Recently this approach was extended to bipedal locomotion [4] with results presented on a simulated SAR-COS robot. The primary contribution was extension of the algorithm to a biped with a full foot plate, which requires a 6 DOF constraint on the foot frame.

Finally, there has also been work in coupling odometry estimates to a higher-level navigation system. Localization of a quadruped (in 3 DOF) against a prior terrain map was explored by [5] using a particle filter.

All of the above works utilize only proprioceptive sensing. While visual mapping is an active area of research — for instance, [6] demonstrated loop-closing and visual keyframe registration to known landmarks in a laboratory setting — it has not been widely adopted for field operations. High computational cost, latency, and sensitivity to environmental

¹Maurice F. Fallon is with the Department of Informatics, University of Edinburgh, UK and the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA maurice.fallon@ed.ac.uk

²Matthew Antone, Nicholas Roy and Seth Teller are with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA antone,nickroy,teller@csail.mit.edu

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE AUG 2014		2. REPORT TYPE		3. DATES COVERED 00-00-2014 to 00-00-2014	
4. TITLE AND SUBTITLE Drift-Free Humanoid State Estimation fusing Kinematic, Inertial and LIDAR sensing				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, 02139				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Under review					
14. ABSTRACT This paper describes an algorithm for the probabilistic fusion of sensor data from a variety of modalities (inertial, kinematic and LIDAR) to produce a single consistent position estimate for a walking humanoid. Of specific interest is our approach for continuous LIDAR-based localization which maintains reliable drift-free alignment to a prior map using a Gaussian Particle Filter. This module can be bootstrapped by constructing the map on-the-fly and performs robustly in a variety of challenging field situations. We also discuss a two-tier estimation hierarchy which preserves registration to this map and other objects in the robot's vicinity while also contributing to direct low-level control of a Boston Dynamics Atlas robot. Extensive experimental demonstrations illustrate how the approach can enable the humanoid to walk over uneven terrain without stopping (for tens of minutes), which would otherwise not be possible. We characterize the performance of the estimator for each sensor modality and discuss the computational requirements.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

conditions (e.g., illumination and visual texture content) all present substantial challenges for robust vision-based mapping onboard a humanoid.

B. Overview

This paper presents a state estimation algorithm that combines measurements from three distinct sensing modalities — inertial, leg kinematics, and LIDAR — into a single consistent estimate of the robot’s pelvis link via probabilistic fusion. The estimator does not require elaborate dynamics models; like [4], we couple foot placements and leg kinematics with inertial predictions in an EKF framework, and like [2], we consider pelvis pose separately from joint states. Our primary novel contributions are (1) incorporation of exteroceptive sensing to achieve reliable drift-free alignment to a prior map while walking; (2) derivation of position corrections from each LIDAR scan via a Gaussian Particle Filter (GPF); and (3) extensive experimental validation of the algorithm on a real humanoid robot.

The estimator is demonstrated using the Boston Dynamics (BDI) Atlas humanoid (Figure 1) provided to our research team for the ongoing DARPA Robotics Challenge (DRC). The LIDAR-based localization component was specifically motivated by the slow locomotion rates achieved in the DRC Trials in December 2013: due to position drift on the order of 2 cm per footstep, teams typically took just two steps at a time to traverse uneven terrain, pausing periodically to manually re-localize the robot and motion plans with respect to the environment. As we push toward enabling greater autonomy in task execution — e.g., walking for several minutes at a time with manipulation actions interspersed — a continuous localization capability becomes critically important, as it allows the robot to retain accurate and consistent reference to terrain maps and objects of interest in its vicinity.

In Section II we present an overview of our requirements for state estimates and discuss two different use cases for our approach. Then in Sections III–V we discuss how each sensor stream can be abstracted to a basic probabilistic measurement suitable for fusion.

Finally in Section VI we benchmark the performance of the algorithm and present results from a series of extended duration walking experiments, most executed passively with BDI’s native controller (totaling approximately one hour), and one directly integrated in the loop with our own controller. A clear operational benefit is demonstrated: alignment to the prior model enables the robot to continuously traverse uneven terrain without stopping, and thus operate up to four times more quickly than previously possible.

II. REQUIREMENTS

The ultimate goal of our research efforts [7] is to develop a system that enables a humanoid robot to operate at a semi-autonomous level with human interaction at a task level, as depicted in Figure 2 : “walk over to the drill and use it to cut a circular hole in the wall”.




Fig. 2. A 3D rendering of the robot with a footstep plan leading towards a fitted model of a (cyan) drill on a (yellow) table.

Executing such actions requires the ability to precisely and continuously localize, thus enabling the robot to walk to a goal without interruption. An alternative would be to encode safety sequences such as stopping short of a goal and then stepping conservatively into position which, as well as being inefficient, adds complexity to an otherwise simple action.

To achieve this, we envisage a navigation system that comprises two types of state estimator, (Figure 3. The first, required by our closed-loop locomotion controller [8], produces reliable estimates of position and velocity with high rate and low latency. Because discontinuities in the position estimate (due to map alignment corrections) or phase errors can cause undesirable control feedback, it is critical that estimates be locally consistent over short time periods.

However, proprioception is subject to drift and thus is not amenable to task-level autonomy, which requires constant long-term localization within the environment but can better tolerate small instantaneous adjustments. We therefore maintain a second localization estimate that incorporates exteroception data from LIDAR (or vision sensors) to reduce drift, but can allow discontinuous corrections of the robot’s position. This two-tiered approach, which is in the spirit of [9], [10], maintains reference to higher-level features of interest used by our robot/operator team. Here we focus primarily on this second, exteroceptive, localization mode; however, in practice we use the same algorithms and software framework for both modes.

III. INTEGRATION

Both estimators utilize the same integrator algorithm, originally described in [11] and used for a micro aerial vehicle (MAV), with simple software configuration flags enabling or disabling the various input sensors.

Following the notation described therein, we wish to estimate the position and orientation of the robot’s kinematic root link, the pelvis, as well as its linear and angular velocities. The full state vector is defined as $x = [w_b^T \ v_b^T \ R \ \Delta^T]^T$ and each component is as follows

- angular velocity, $w_b = [p \ q \ r]^T$

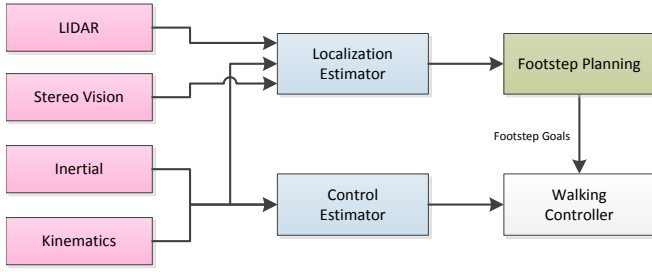


Fig. 3. To allow possibly discontinuous exteroception corrections we propose a hierarchy in which these corrections are fed to the walking controller indirectly and at low rate (in this case as footstep goals via our planner). Note that the integration of stereo vision is ongoing work.

- linear velocity, $v_b = [u \ v \ w]^T$
- orientation, R
- position, $\Delta = [\Delta_x \ \Delta_y \ \Delta_z]^T$

Both velocity components are expressed in the pelvis frame, while the position and orientation of the pelvis are expressed in a fixed world frame. We exclude both the robot’s contact foot position and the joint states from this state vector and filter them separately (as do others including [2]).

The pelvis of the Atlas humanoid contains the robot’s primary Inertial Measurement Unit (IMU), located 9 cm behind the pelvis link position and rotated by 45° . The sensor is a KVH 1750-IMU comprised of Fiber Optic Gyroscopes (FOG) and Micro Electromechanical Systems (MEMS) accelerometers of tactical grade. The IMU thus provides excellent performance compared to sensors used by other legged robots.

The state estimate and its associated covariance are updated using an Extended Kalman Filter (EKF). The prior distribution is propagated using a process model driven by the IMU measurements: the rotation rate w_b and the acceleration a_b are both sensed in the IMU frame and transformed to the body frame¹. As discussed in [11], orientation uncertainty is expressed in exponential coordinates around the body frame.

The IMU sensor provides very accurate raw measurements, but vibrations induced by the hydraulic pressurizer within the robot’s torso corrupt the signal. We therefore apply a cascading set of IIR-notch filters to dampen the 85 Hz vibrational component and its harmonics.

Additional states of the EKF are used to maintain rotation rate and acceleration bias estimates, which are computed while the robot stands still at the start of an experiment. Although the estimator supports on-line bias updates, we typically retain these initial values as they tend to remain consistent throughout. Via a static analysis, we also determined a characteristic scaling bias of approximately 0.04 m/sec^2 in the accelerometer’s Z-axis.

In the subsequent sections we will describe how each individual sensing modality is used to form Kalman measurement updates to the state vector. This information is summarized in Table I.

¹The manufacturer-provided orientation estimate was not used but is compared with in Section VI.

Dimension	Pos Δ	Orient R	Velocity v_b	Ang Rate w_b	Accel a_b
Accelerometers					✓
Gyroscopes				✓	
Leg sensing	✗	✗	✓	✗	
LIDAR	✓	✓			
Visual Odometry	✓	✓			

TABLE I

CONTRIBUTION OF VARIOUS SENSORS TO THE FILTERED STATE ESTIMATE. MODES OF INTEGRATION FOUND TO BE USEFUL ARE MARKED ✓ AND THOSE NOT USED (FOR A VARIETY OF REASONS) ARE INDICATED ✗.

IV. LEG KINEMATICS

The robot has two legs, each with six joints: three at the hip, a knee joint and two angle joints. As with many leg kinematic integration algorithms [1], [3], our approach assumes that the robot’s stance foot maintains non-slipping contact with the ground during part of the gait and that this foot is stationary. This allows instantaneous velocity and position measurements of the robot’s pelvis to be inferred via forward kinematics. Of course in practice perfectly clean contacts and stable feet are seldom achieved, but we assert that for short periods (the sample time of our sensors) these assumptions are reasonable.

On the Atlas robot, the angle of each leg joint is sensed by measuring the travel of its hydraulic actuator using an LVDT and then computing a transform through the joint linkage. This model does not account for flexion when loaded or stiction when a joint changes direction. The robot’s arm joints, in contrast, contain additional encoders that directly sense the joint angle, which we deem to be much more accurate. Figure 4 shows a comparison between LVDT and encoder values measured while continuously actuating the left shoulder joint, with differences of a degree or more. Because of these un-sensed effects, leg kinematic integration is subject to position drift at rates that vary due to factors such as dynamism of the walking gait and controller execution.

A. Contact Classification

At the base of the locomotion algorithm, a gait transition detector infers the current stage of the walking motion and then decides which of the feet has stationary contact with the ground. We use a Schmitt trigger with a threshold of 575 N to classify contact forces sensed by the robot’s 3-axis foot force-torque sensors and detect whether either foot is in contact (these sensors are calibrated at the start of experiments while the robot is suspended above the ground). A simple state machine then decides which foot has the most reliable contact and thus will provide the basis for kinematic measurements. In the specific case of walking up stairs, the toe of the trailing foot can be used to push the robot’s pelvis forward and upward while not being in stationary contact (known as “toe off”). In this situation we ensure that the leading foot is assigned to be the primary fixed foot.

We also classify other events in the gait cycle. Striking contact is determined when force of 20–30 N is maintained




Fig. 4. Comparison between LVDT potentiometer and direct encoder joint sensing for the robot's left shoulder joint. The joint was actuated between its limits in each direction 3 times illustrating non-linear directional effects. Such effects are not sensed by the leg joint sensors.

for more than 5 msec. Breaking contact is determined when force falls below 275 N in the opposite direction. Because these events create sensing spikes, we apply higher measurement covariance in the EKF when they occur.

We note that when the robot is in a double support stance, information from both legs can be leveraged to provide additional kinematic measurements. For simplicity we currently neglect this information.

B. Kinematic Measurements

Given the high accuracy of IMU orientation data, we choose to use joint sensing to measure only the linear position and velocity of the pelvis. Thus, our initial assumptions are that the pelvis orientation R_t^p is known and that a stationary foot position $\Delta_t^f = [\Delta_{x,t}^f, \Delta_{y,t}^f, \Delta_{z,t}^f]^T$ has been decided upon. After pre-filtering the incoming joint angles with a low-pass 16 Hz FIR filter, we solve for forward kinematics between the pelvis and the foot. With these two constraints we can then recover the inferred orientation of the foot at the current time R_t^f and a nominal position of the pelvis at that time Δ_t^p .

Two types of filter measurement can be extracted from this calculation. The first approach we considered was to estimate the foot position Δ_t^f for the duration of a stride, which in turn would result in a measurement of the pelvis position. However, because of the inconsistencies in joint sensing and because the robot's foot does not make and maintain perfectly clean static contact with the ground, an alternative is to instead use the difference between consecutive position estimates over a short period of time to create a velocity measurement:

$$\hat{v}_b^m = \frac{\Delta_t^f - \Delta_{t-1}^f}{T_s} \quad (1)$$

where T_s is the integration period, typically 3 msec.

This approach is attractive because each resultant observation is a discrete measurement of the robot's velocity and




Fig. 5. Top: Evolution of foot force signals for the left (green) and right (red) leg during two steps. Bar Chart: classification of the primary standing foot (light color indicates a ground contact event). Lower plots: pelvis lateral (body Y-direction) velocity estimates for (1) raw leg odometry (the estimate during contact events is shown in red), (2) filter output and (3) VICON-based ground truth.

does not retain the accumulated effect of non-ideal ground contact, e.g., the footplate rolling or sliding. The influence of an erroneous velocity is transient and quickly corrected by subsequent observations.

Using both measurement types together would be comparable to the approach in [2], but we avoid doing so as it raises the possibility of creating inconsistencies, particularly when combined with position measurements derived from the LIDAR module (presented in the following section).

We note that our approach neither assumes knowledge of the terrain surface normal that robot is standing upon nor attempts to maintain a consistent estimate of the foot orientation over time.

Figure 5 contains a number of plots comparing (1) the raw pelvis velocity measurements inferred from kinematics with (2) the output of our integrating filter and (3) the velocity estimated from VICON motion capture. Typical pelvis velocity standard deviations, measured when standing still, are as follows:

- 1) Raw incoming kinematics: 7.6 cm/sec

- 2) After joint level filtering: 2.3 cm/sec
- 3) After EKF integration: 1.4 cm/sec

V. LIDAR MEASUREMENTS

As illustrated in Figure 1, the robot is equipped with a Multisense SL sensor head designed by Carnegie Robotics which combines a fixed binocular stereo camera with a Hokuyo UTM-30LX-EW planar LIDAR sensor mounted on a spindle that can rotate up to 30 RPM. (We typically operate the device at 5 RPM to densely sample the terrain when walking). The LIDAR captures 40 scan lines of the environment per second, each containing 1081 range returns out to a maximum range of 30 meters. The entire head can pitch up and down (powered by a hydraulic actuator) but cannot yaw or roll.

Our projection of LIDAR range returns as points in the 3D workspace explicitly accounts for the rotation of the LIDAR’s internal mirror during its 1/40 second scanning period. Neglecting this effect results in mis-projections of lateral returns by as much as 2.5 m at the highest spindle speed. Accurate projection also requires precise calibration of the LIDAR sensor, performed offline as discussed in [7].

A. Contribution to Estimation

Our strategy is to use the LIDAR to *continuously* infer the robot’s position relative to a prior map while walking. We cannot assume that the sensor is aligned horizontally [12], nor can we afford time to stop moving and perform static 3D registration, e.g., using an Iterative Closest Point algorithm [13]. Instead we aim to incorporate information from each individual LIDAR scan into the state estimate using a Gaussian Particle Filter, as originally described in [11].

In typical operation, the robot is first commanded to stand still for between 10 and 30 seconds while it collects a full 3D point cloud of its environment (see Figure 6). This cloud is then converted into a probabilistic occupancy grid (OctoMap [14]) against which efficient localization comparisons are later performed. While the MAV experiments presented in [11] required offline mapping with a separate sensor, our legged humanoid and actuated LIDAR with 30 m range permit the map to be constructed immediately prior to operation and utilized during the entire task². This makes our approach practical both for highly variable laboratory experiments and for field trials in which the environment is initially unknown. Furthermore, if the robot were to approach the map boundary, on-line construction of a new map could easily be performed by standing still again.

Since the LIDAR is fundamentally a planar 2D sensor, only a subset of the state vector (namely x , y and yaw in the current sensor plane) is observable at any given instant. We therefore partition the full state vector into observable and unobservable components, and use a GPF to incorporate each laser measurement over the observable variables. The particle filter samples are weighted according to the proposed sub-state likelihood, which is computed by comparing the LIDAR




Fig. 6. The robot initially collects a static LIDAR point cloud of its environment, which is then converted into an occupancy map for subsequent localization.

measurements, projected from the sub-state, to the prior map. From these weighted samples a mean and covariance, and in turn an equivalent Kalman measurement update for the full state vector, are calculated.

While this formulation provides some information about the full pose of the robot, the inertial measurements of pitch and roll are already highly accurate. Therefore, to decrease computation and increase robustness, we incorporate only the position Δ and the yaw in the particle filter step rather than all observable variables.

B. Latency and Computation

Typical latencies of the kinematic and LIDAR components are as follows, for a 3.3GHz 12 core desktop PC:

Component	Latency	Frequency
Lower joint Kalman Filters	0.16 msec	1 kHz
Pose Extended Kalman Filter	0.54 msec	333 Hz
LIDAR data transmission	7 msec	40 Hz
GPF processing time	11.4 msec	40 Hz
Overall LIDAR latency	18.4 msec	40 Hz

These values and the experiments presented in Section VI use 1000 GPF samples, although reliable performance (and reduced latency) is possible with just 300 samples.

We use a multi-process messaging architecture to parallelize computation, with the GPF algorithm requiring a single CPU core. Within the estimator, the EKF retains a 1 sec history of measurements to accommodate the LIDAR and GPF latencies. Corrections are applied out-of-order to ensure timely estimates.

C. Reliability and Practicality

We also considered the reliability of each modality within the DRC competition context. A practical issue with the inertial sensor is that the robot must be completely stationary during initialization. Boston Dynamics’ estimator requires initialization as soon as the robot is first powered on, typically with its feet solidly contacting the ground. Our proposed estimator provides greater flexibility, as it can

²The DRC Trials terrain course was approximately 15 m in length.

be initialized at any point when the robot is standing and deemed to be stationary.

While the time taken to construct the LIDAR map (10–30 sec) is a minor inconvenience, in all datasets and task scenarios available to us there was sufficient stationary structure within sensor range for localization to operate reliably. Since the algorithm uses measurements of the entire environment, movement of a few objects or people near the robot has no noticeable effect on performance. However, in scenarios containing substantial background motion (such as the crowds attending the DRC Trials), special consideration may be required to disregard portions of the map with significant activity.

Finally, while our preference has been to build a map from a single location and use it continuously while operating, in Section VI-B we demonstrate that the GPF localization algorithm is robust to non-encompassing maps as well.

VI. EXPERIMENTS

In this section we demonstrate the performance of the state estimator through a variety of experiments.

In each case our team’s footstep planner (described in [7] and further developed in [15]) creates a kinematically-feasible footstep sequence that reaches a goal position and orientation while minimizing the number of steps taken. The experiments involving uneven terrain traversal require human input to ensure correct initial footstep placement.

In our first set of experiments, the system sequentially feeds footsteps to the Boston Dynamics walking controller (with its own internal state estimate) for execution³ while continually modifying their positions during execution so that the controller achieves the intended motion. Figure 7 summarizes our results for a variety of walking patterns totaling 57 minutes of operation and 155 meters traveled.

The kinematic-only estimates (both our own and the manufacturer’s) are seen to continuously drift, typically at 1.2–1.5 cm per step. This drift rate generally disimproves when the locomotion is atypical: up inclines, dynamic, or with extended steps. Orientation estimation is comparable; however, when the robot translates by any significant amount, the value of fusing LIDAR-based corrections quickly becomes evident. In the manipulation experiment, the LIDAR contribution actually degrades performance slightly; for this reason, and for general stability, we now discard LIDAR data when standing still.

A. Continuous Terrain Traversal

In one specific experiment of note, depicted in Figure 8, the robot traverses a set of cinder blocks that matches the terrain course from the December 2013 DRC Trials. The LIDAR-aided localization estimate remains accurate to within 2 cm at all times, enabling the robot to walk continuously. It can even execute this sequence in reverse to return to its exact starting position, despite having no rear-facing sensors, because the forward-facing sensors keep the

³The controller provides both a quasi-static step mode and a dynamic walking mode; we almost exclusively use the former.

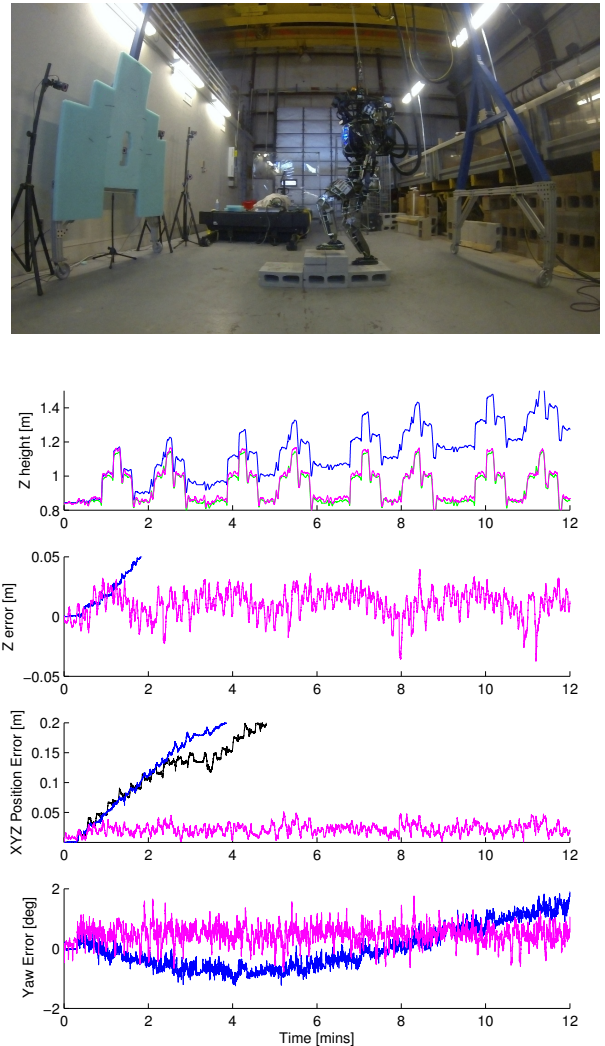


Fig. 8. Analysis of continuous traversal of the terrain illustrated in the upper image. In the upper plot, green is the VICON ground truth; blue is the estimate provided by BDI; black incorporates inertial and kinematic data and behaves similarly; magenta additionally incorporates LIDAR and as a result remains accurately localized. A longer 5-block traversal sequence (both forward and backward) can be seen in the supplementary video accompanying this paper.

robot precisely localized. This forward-backward traversal was repeated 4 times in 12 minutes.

By comparison, the kinematic-only estimates drift continuously, culminating in a total of 0.8 m accumulated error, and would have caused the robot to fall after its first few steps. During the December 2013 Trials, most teams executed the terrain course two steps at a time because of this drift. Extrapolating from this experiment, the proposed algorithm can enable execution up to four times more quickly than was achieved in the Trials.

B. Localization in Partial Maps

In a second noteworthy experiment, captured in a video accompanying this paper, we explored the performance of the estimator when the robot is not fully surrounded by

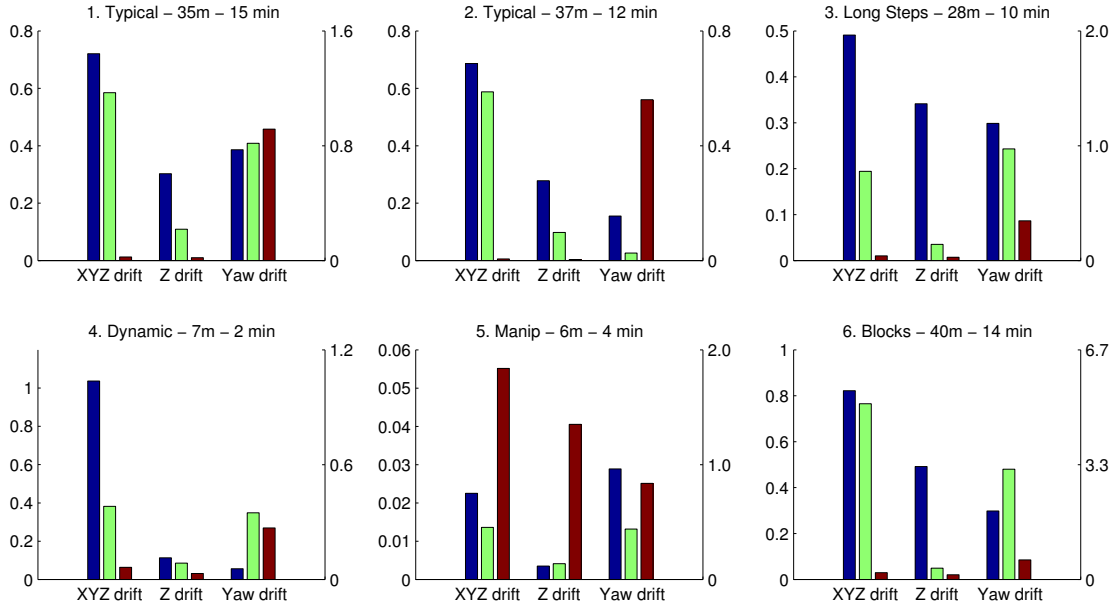


Fig. 7. Summary of localization accuracy for a variety of walking experiments. Position drift is measured with the left scale (in meters) and yaw drift with the right scale (in degrees). Error (versus ground truth) of the BDI state estimator (blue), MIT’s kinematic-only (green) and closed-loop LIDAR (magenta) estimators are shown. Clockwise from upper left: 1: typical gait (15 cm forward steps), 2: typical gait with a partial map (see Section VI-B, 3: long steps (36 cm forward steps), 4: dynamic walking, 5: carrying out manipulation, 6: traversing the cinder block course in Figure 8.

the prior map. The field of view of the LIDAR sensor roughly accommodates mapping the hemisphere in front of the robot from a single standing position. In this experiment the robot first created a map from its starting pose and then turned 90 degrees before walking for 12 minutes, such that a significant portion of the LIDAR data fell outside of the map. Localization performance (Experiment 2 in Figure 7) was much the same as in the other experiments, further demonstrating the robustness of this approach.

The likelihood function underlying the GPF is computed as the product of the likelihood of each LIDAR return in a scan. When a return falls far from an occupied cell of the OctoMap, an unobserved likelihood is applied to this return, resulting in stable performance despite a high percentage of potentially corruptive outlier measurements.

VII. CONCLUSIONS

We have presented a probabilistic fusion algorithm for humanoid state estimation and characterized the contribution of inertial, kinematic and LIDAR sensing to that estimate. Of particular note is that the approach supports continuous *drift free* localization within a 3D LIDAR map, which we aim to use for enabling longer-term semi-autonomous operation of the BDI Atlas humanoid robot. A key demonstration is that our approach allows the robot to continuously walk over complex terrain while precisely achieving all of its footstep placements, which was demonstrated in a variety of extended experiments and would not otherwise have been possible.

In future work we will focus on better characterization of the joint angle biases so as to reduce the rate of kinematic

drift as well as incorporating vision as a contributing component (Section VII-A). Meanwhile, the MIT DRC team’s walking controller currently under development is using this estimator as its primary input (Section VII-B). A selection of videos of the algorithm operating, including state estimation combined with walking and balancing control, can be seen at people.csail.mit.edu/mfallon/state-est

A. Stereo Vision

Several obstacles impede the adoption of vision in practical outdoor scenarios. Figure 10 illustrates a challenging but realistic scene from the DRC Trials containing strong shadows, false feature detections, and abrupt lighting changes, all of which are detrimental to vision algorithms. Computational requirements and latency are also not insignificant; we have therefore focused efforts to date on exteroceptive measurements from LIDAR alone. Nonetheless, we have also experimented with stereoscopic sensing as a potential complementary modality. The entire state space is observable, since each measurement covers a full volume of space rather than a single planar slice; and the output of visual odometry algorithms can be formulated suitably for incorporation into our probabilistic fusion framework.

The Multisense SL’s stereo camera captures 1-MPix images at up to 15 Hz with an 80° x 80° field of view. An onboard FPGA rectifies the images and calculates a dense disparity map (which can be converted to a depth map with a baseline of 7 cm) at frame rate using Semi-Global Matching [16]. We benchmarked the performance of FOVIS [17], a state-of-the-art visual odometry system, by post-processing




Fig. 9. Overhead view of open-loop visual odometry (magenta) and VICON-based ground truth (green)—both estimating the position of the pelvis link—for a typical walking gait. The robot starts at (0, 0.4) and faces predominantly to the right throughout. After 40 steps, the position estimate drifted by only 3 cm. Note that the robot was outside of the VICON field of view for a short time in the upper right.



Fig. 10. In realistic field scenarios such as the DRC Trials, strong shadows and self observations result in false visual features (shown in red).

color and disparity images generated at 10Hz during a typical walking experiment (Experiment 1 in Figure 7) to produced the plot in Figure 9. In a laboratory setting and with a smooth walking motion, FOVIS produces negligible drift and retains alignment to key-frames for several steps at a time, with ~ 150 msec typical latency.

B. Adaptations for Closed Loop Control

In parallel with these efforts, we have been adapting our quadratic program-based walking controller [8] for use in place of the manufacturer-provided controller. The state estimator described here is being used directly in the control loop. In addition to the pelvis state estimation, we also carry out independent Kalman filtering of the 12 lower body joints to produce joint positions and velocities. Further filtering of the pelvis rotation rates has been useful to minimize control feedback. As yet, we have not moved to integrate the LIDAR-based position correction module into our real-time control.

ACKNOWLEDGMENT

This work was supported by the Defense Advanced Research Projects Agency (via Air Force Research Laboratory

award FA8750-12-1-0321). We thank the MIT DRC Team for their work on the humanoid system and Adam Bry and Abe Bachrach for development of the core estimation framework.

REFERENCES

- [1] B. J. Stephens, "State estimation for force-controlled humanoid balance using simple models in the presence of modeling error," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, May 2011, pp. 3994–3999.
- [2] X. Xinjilefu, S. Feng, W. Huang, and C. Atkeson, "Decoupled state estimation for humanoids using full-body dynamics," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Hong Kong, China, 2014.
- [3] M. Bloesch, M. Hutter, M. A. Hoepfner, S. Leutenegger, C. Gehring, C. D. Remy, and R. Siegwart, "State estimation for legged robots - consistent fusion of leg kinematics and IMU," in *Robotics: Science and Systems (RSS)*, 2012.
- [4] N. Rotella, M. Bloesch, L. Righetti, and S. Schaal, "State estimation for a humanoid robot," *CoRR*, vol. 1402.5450, 2014, IROS Submission.
- [5] S. Chitta, P. Vernaza, R. Geykhman, and D. Lee, "Proprioceptive localization for a quadrupedal robot on known terrain," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, April 2007, pp. 4582–4587.
- [6] O. Stasse, A. J. Davison, R. Sellaoui, and K. Yokoi, "Real-time 3D SLAM for a humanoid robot considering pattern generator information," in *IEEE/RSS Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2006.
- [7] M. Fallon, S. Kuindersma, S. Karumanchi, M. Antone, T. Schneider, H. Dai, C. P. D'Arpino, R. Deits, M. DiCicco, D. Fourie, T. T. Koolen, P. Marion, M. Posa, A. Valenzuela, K.-T. Yu, J. Shah, K. Iagnemma, R. Tedrake, and S. Teller, "An architecture for online affordance-based perception and whole-body planning," *J. of Field Robotics*, 2014, under review.
- [8] S. Kuindersma, F. Permenter, and R. Tedrake, "An efficiently solvable quadratic program for stabilizing dynamic locomotion," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Hong Kong, China, June 2014.
- [9] D. C. Moore, A. S. Huang, M. Walter, E. Olson, L. Fletcher, J. Leonard, and S. Teller, "Simultaneous local and global state estimation for robotic navigation," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2009.
- [10] E. Marder-Eppstein, E. Berger, T. Foote, B. Gerkey, and K. Konolige, "The Office Marathon: Robust navigation in an indoor office environment," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, May 2010.
- [11] A. Bry, A. Bachrach, and N. Roy, "State estimation for aggressive flight in GPS-denied environments using onboard sensing," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2012, pp. 1–8.
- [12] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte Carlo localization for mobile robots," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, May 1999.
- [13] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, no. 2, pp. 239–256, 1992.
- [14] K. M. Wurm, A. Hornung, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: A probabilistic, flexible, and compact 3D map representation for robotic systems," in *Proc. of the ICRA 2010 Workshop on Best Practice in 3D Perception and Modeling for Mobile Manipulation*, Anchorage, AK, USA, May 2010.
- [15] R. L. Deits and R. Tedrake, "Computing large convex regions of obstacle-free space through semi-definite programming," in *Workshop on the Algorithmic Fundamentals of Robotics (WAFR)*, Istanbul, Turkey, Aug. 2014.
- [16] H. Hirschmüller, "Accurate and efficient stereo processing by semi-global matching and mutual information," in *Proc. IEEE Intl. Conf. Computer Vision and Pattern Recognition*, vol. 2, June 2005, pp. 807–814.
- [17] A. Huang, A. Bachrach, P. Henry, M. Krainin, D. Maturana, D. Fox, and N. Roy, "Visual odometry and mapping for autonomous flight using an RGB-D camera," in *Proc. of the Intl. Symp. of Robotics Research (ISRR)*, Flagstaff, USA, Aug. 2011.